

---

# Can Pretrained Neural Networks Detect Anatomy?

---

**Vlado Menkovski**  
Philips Research  
Eindhoven, Netherlands  
vlado.menkovski@philips.com

**Zharko Aleksovski**  
Philips Research  
Eindhoven, Netherlands  
zharko.aleksovski@philips.com

**Axel Saalbach**  
Philips Research  
Hamburg, Germany  
axel.saalbach@philips.com

**Hannes Nickisch**  
Philips Research  
Hamburg, Germany  
hannes.nickisch@philips.com

## Abstract

Convolutional neural networks demonstrated outstanding empirical results in computer vision and speech recognition tasks where labeled training data is abundant. In medical imaging, there is a huge variety of possible imaging modalities and contrasts, where annotated data is usually very scarce. We present two approaches to deal with this challenge. A network pretrained in a different domain with abundant data is used as a feature extractor, while a subsequent classifier is trained on a small target dataset; and a deep architecture trained with heavy augmentation and equipped with sophisticated regularization methods. We test the approaches on a corpus of X-ray images to design an anatomy detection system.

## 1 Introduction and Motivation

Deep learning [9] (DL) is rapidly gaining momentum in computer vision [5], speech recognition [3] and medical imaging applications due to their record-breaking empirical performance and their promise to jointly learn the low-level features and the high-level decisions directly from raw data.

One major open research challenge is the transfer of the excellent predictive performance of DL methods, in particular convolutional neural networks (CNN), from the “data-laden” regime of computer vision<sup>1</sup> to the “data-bounded” regime of medical imaging.

Since the typical number of parameters in the DL models is very big the major concern for operating in “data-bounded” regimes is inability to generalize and form robust features because the small number of examples do not capture well the variability in the input space. There are various approaches to deal with this challenge. Data augmentation can be used to increase the effective size of the training set and drive the training to build invariances to specific transformations. Network architecture design can be used to hard-code prior knowledge (such as translation invariance) by design. Sophisticated regularization methods such as dropout [8] and batch normalization (BN) [4] can work to counter overfitting. Finally, *global* training of the feature extraction layers on abundant data domains (i.e. natural images) and then training the “classification” higher layers from the limited data domains can be used to ensure that the lower-level features are invariant to the particularities of the training dataset.

We choose two directions to investigate these approaches for the goal of anatomy detection in X-ray images. The first using a network trained on a large corpus of natural images for feature extraction

---

<sup>1</sup>There are  $10^6$  images in  $10^3$  classes in the ILSVRC challenge, <http://www.image-net.org/challenges/LSVRC>.

combined with a SVM [1] classifier trained on the medical images. And a second approach using heavy augmentation and regularization of a deep neural network trained only on the small-sized medical image dataset.

## 2 Experiments and Results

In order to assess the performance of specifically trained and off-the-shelf networks, radiographs from the ImageClef 2009 - Medical Image Annotation task<sup>2</sup> were used. This challenge database consists of a broad range of X-ray images from clinical routine, with a detailed anatomical classification. For convenience, we flattened the hierarchical class representation and disregarded classes with fewer than 50 example images. In doing so we ended up with 24 unique classes. For evaluation purposes, the entire data corpus (of 14676 images) was divided into a training and test set, covering 90% and 10% of the data respectively.

As a first baseline, the OverFeat network (see [6]) in combination with a (linear) multi-class SVM was employed. OverFeat is a convolutional neural network that was trained on 1.2 million non-medical images from the ImageNet2012 database with 1000 classes, and has been successfully employed as a generic feature extractor in different applications.

For the purpose of this experiment, we selected the *fast* OverFeat instantiation and used the default setting for feature extraction. This results in a 4096 dimensional representation of the images, which was subsequently used as input for the SVM classifier. The regularization parameter C of the linear SVM was estimated on the training set using exhaustive search and 5-fold cross-validation.

For the second direction we designed a network 2 inspired by the work of Razavian et al. [7], using dropout and batch normalization with leaky ReLU activation functions [2]. In this case for training we used images of 128 by 128 pixels with various augmentations ranging from resizing and cropping, rotation, translation, shearing, stretching and flipping. The results of the two directions are given in table 1.

Table 1: Results: Accuracy in percent

OverFeat + SVM	Custom CNN
92.42	95.14

## 3 Conclusions and Perspectives

We investigated when, where, and how to (re)train a convolutional neural network used for anatomy detection. Our results suggest that pretrained networks are good image descriptors even outside their training domain and that retraining only the last layer is a viable alternative to perform transfer learning in light of limited training data. We also demonstrated that sufficiently expressive models such as a very deep neural network can be trained even on relatively small number of annotations if proper augmentation and regularization is implemented. Our work is a first step towards a multi-modality anatomical expert system with components trained to maximize data efficiency.

## References

- [1] C. Cortes and V. Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.
- [2] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *arXiv preprint arXiv:1502.01852*, 2015.
- [3] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *Signal Processing Magazine, IEEE*, 29(6):82–97, 2012.

<sup>2</sup><http://www.imageclef.org/2009/medanno>

- [4] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [5] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems 25 (NIPS)*, 2012.
- [6] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. In *International Conference on Learning Representations (ICLR 2014)*, 2014.
- [7] A. Sharif Razavian, J. Sullivan, A. Maki, and S. Carlsson. A baseline for visual instance retrieval with deep convolutional networks. In *International Conference on Learning Representations, May 7-9, 2015, San Diego, CA. ICLR, 2015*.
- [8] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [9] Y. B. Yann LeCun and G. Hinton. Deep learning. *Nature*, 521:436–444, 2015.

## Annex

Table 2: Custom Deep Neural Network architecture

ConvLayer + BN	(3x3, 32x)
ConvLayer + BN	(3x3, 16x)
MaxPoolLayer	(3x3, 2x2 stride)
ConvLayer + BN	(3x3, 64x)
ConvLayer + BN	(3x3, 32x)
MaxPoolLayer	(3x3, 2x2 stride)
ConvLayer + BN	(3x3, 128)
ConvLayer + BN	(3x3, 128)
ConvLayer + BN	(3x3, 64)
MaxPoolLayer	(3x3, 2x2 stride)
ConvLayer + BN	(3x3, 256)
ConvLayer + BN	(3x3, 256)
ConvLayer + BN	(3x3, 128)
MaxPoolLayer + Dropout	(3x3, 2x2 stride)
DenseLayer + BN + Dropout	(256)
DenseLayer + BN + Dropout	(256)
Softmax	(24)